

A Meta-Programming Approach to Realizing Dependently Typed Logic Programming

Zachary Snow

Computer Science and Engineering
University of Minnesota
200 Union Street SE
Minneapolis, MN 55455
snow@cs.umn.edu

David Baelde

Computer Science and Engineering
University of Minnesota
200 Union Street SE
Minneapolis, MN 55455
dbaelde@cs.umn.edu

Gopalan Nadathur

Computer Science and Engineering
University of Minnesota
200 Union Street SE
Minneapolis, MN 55455
gopalan@cs.umn.edu

Abstract

Dependently typed λ -calculi such as the Logical Framework (LF) can encode relationships between terms in types and can naturally capture correspondences between formulas and their proofs. Such calculi can also be given a logic programming interpretation: the Twelf system is based on such an interpretation of LF. We consider here whether a conventional logic programming language can provide the benefits of a Twelf-like system for encoding type and proof-and-formula dependencies. In particular, we present a simple mapping from LF specifications to a set of formulas in the higher-order hereditary Harrop (*hohh*) language, that relates derivations and proof-search between the two frameworks. We then show that this encoding can be improved by exploiting knowledge of the well-formedness of the original LF specifications to elide much redundant type-checking information. The resulting logic program has a structure that closely resembles the original specification, thereby allowing LF specifications to be viewed as *hohh* meta-programs. Using the Teyjus implementation of λ Prolog, we show that our translation provides an efficient means for executing LF specifications, complementing the ability that the Twelf system provides for reasoning about them.

Categories and Subject Descriptors D.3.2 [Programming Languages]: Language Classifications— Constraint and logic languages; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic— Lambda calculus and related systems, Logic and constraint programming, Proof theory

General Terms Theory, Languages

Keywords logical frameworks, dependently typed lambda calculi, higher-order logic programming, translation

1. Introduction

There is a significant, and growing interest in mechanisms for specifying, prototyping and reasoning about formal systems that are described by syntax-directed rules. Dependently typed λ -calculi such as the Logical Framework (LF) [11] provide many conveniences

from a specification perspective in this context. Such calculi facilitate the use of a higher-order approach to describing the syntax of formal objects and they allow relationships between terms to be captured in an elegant way through type dependencies. Furthermore, dependently typed λ -calculi enjoy a well-known isomorphism between formulas and types [12], leading to a unification of the concept of a proof of a formula with an inhabitant of a given type. Thus, the search for type inhabitants can be identified with proof-search and can thereby be given a logic programming interpretation. The Twelf system [19] that we consider here exploits these possibilities relative to LF. As such, it has been used successfully in specifying and prototyping varied formal systems, and mechanisms have also been built into it to reason about specifications.

Predicate logics are also capable of encoding syntax-directed specifications, and provide the basis for logic programming languages in the familiar tradition of Prolog. Within this framework, the logic of higher-order hereditary Harrop (*hohh*) formulas [15] that underlies the language λ Prolog [16] provides a builtin ability to treat binding notions in syntax and thus has particular usefulness in representing formal systems. However, unlike LF, this logic cannot reflect dependencies between objects into types and does not directly represent the relationship between formulas and their proofs. While such correspondences can always be encoded by hand through auxiliary predicate definitions, it is of interest to understand if a systematic encoding is possible. A specific form to this question is if Twelf specifications can be translated into λ Prolog programs, allowing such specifications to be seen as λ Prolog “meta-programs.” There are benefits to such a possibility: the convenience of writing specifications using dependent types can be combined with the ability both to execute them via an efficient λ Prolog implementation, and to reason about them using logics and systems meant for analyzing *hohh* descriptions [2, 7, 10, 14].

A partial answer to the question raised above has been provided by Felty, who described a translation of LF specifications to *hohh* formulas and then showed that LF derivations correspond exactly to *hohh* derivations of the translated LF judgment [5, 6]. The focus on matching derivations allows Felty to assume the existence of a complete LF judgment, and, in particular, of an LF object in her translation. However, this assumption is inappropriate in our context, given that we are interested in *constructing* proof terms that show particular types are inhabited, *i.e.*, in *proof search* that plays a fundamental role in the logic programming setting. We therefore refine the earlier mapping to remove this assumption and show that the resulting translation preserves derivability in a sense relevant to the logic programming interpretation; an important part of our proof is showing how to extract an LF object satisfying

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PPDP'10, July 26–28, 2010, Hagenberg, Austria.

Copyright © 2010 ACM 978-1-4503-0132-9/10/07...\$10.00

a type from a derivation constructed using the *hohh* version of the specification. Our first encoding may include redundant type-checking judgments which obscure the translated specification and can result in poor execution behavior. We design conditions for eliminating some of these judgments, resulting in an improved translation that corresponds closely to the intention of the original LF specification. This part of our work relies on an analysis of the structure of LF expressions and also has relevance, for example, to providing compact representations of proof terms. Finally, we demonstrate that the execution of the translated form by means of the Teyjus implementation [9] of λProlog [16] provides an effective means for animating Twelf programs.

In the next two sections, we describe a relevant fragment of the *hohh* logic and the Twelf specification language. Section 4 then presents our first translation. In the following section, we describe and exploit a property of LF expressions and type-checking to refine the earlier translation, producing a more efficient and transparent version. Section 6 provides experimental data towards supporting the use of this translation as a means for executing Twelf programs. We conclude the paper with a discussion of related work and possible future directions. This work has been developed in [24]; we refer the reader to that document for complete proofs and more detailed discussions.

2. A Higher-Order Predicate Logic for Describing Computations

The logic of *hohh* formulas is based on an intuitionistic version of Church’s simple theory of types [4]. Both logics are built over a typed form of the λ-calculus. The types are constructed using \rightarrow , the infix, right associative function type constructor, starting from a finite collection of atomic types that includes *o*, the type of propositions, and at least one other type.¹ We assume that we are given sets of variables and constants, each with an associated type. The full collection of (typed) terms is generated from these by the usual abstraction and (left associative) application operators. Terms that differ only in the names of their bound variables are not distinguished. We further assume a notion of equality between terms that is generated by β - and η -reduction. It is well-known that every term has a unique normal form under these reduction operations in this simply-typed setting. All terms are to be converted into such a form prior to their consideration in any context. We write $t[s_1/x_1, \dots, s_n/x_n]$ to denote the result of simultaneously replacing the variables x_1, \dots, x_n with the terms s_1, \dots, s_n in the term t , renaming bound variables as needed to avoid accidental capture. This substitution operation is defined only when s_i and x_i are of the same type for $1 \leq i \leq n$.

We will use only a fragment of the full *hohh* logic here; this fragment still possesses the proof-theoretic properties that are fundamental to the logic programming interpretation of the *hohh* logic. The constants from which terms are constructed are differentiated into *nonlogical* ones that constitute a *signature* and logical ones. We do not permit *o* to appear in the type of the arguments of non-logical constants and variables. The logical constants are restricted to \top of type *o*, \supset of type $o \rightarrow o \rightarrow o$ that is written in the customary infix form and, for each type α , Π of type $(\alpha \rightarrow o) \rightarrow o$. Π represents the universal quantifier as a function over sets. We abbreviate $\Pi (\lambda x.F)$ by $\forall x.F$. An *atomic formula*, denoted by A , is a term of type *o* of the form $p \ t_1 \ \dots \ t_n$ where p is a nonlogical constant. The logic of interest is characterized by two collections of terms called *G*- and *D*-formulas that are defined mutually recursively by the following syntax rules:

$$\begin{aligned} G &:= \top \mid A \mid D \supset G \mid \forall x.G \\ D &:= A \mid G \supset D \mid \forall x.D \end{aligned}$$

A specification or logic program is a finite collection of closed *D*-formulas that are also called *program clauses* and a goal or a query is a closed *G*-formula.

Computation corresponds to searching for a derivation of a sequent of the form $\Sigma; \Gamma \longrightarrow G$ where Σ is the initial (language) signature, Γ is a logic program and G is a goal. Figure 1 presents the rules for constructing such a derivation. Read in a proof search direction, the $\forall R$ rule leads to an expansion of the signature in the sequent whose derivation is sought and the $\supset R$ rule similarly causes an addition to the logic program. The expression “ t is a Σ -term” in the $\forall L$ rule means that t is a closed term all of whose nonlogical constants are contained in Σ . The derivation rules manifest a goal-directed character: to find a derivation for $\Sigma; \Gamma \longrightarrow G$, we simplify G based on its logical structure and then use the *decide* rule to select a formula from the logic program for solving an atomic goal. Notice also that the *decide* rule initiates the consideration of a focused sequence of rules that is similar to backchaining.² In particular, if the formula selected from Γ has the structure

$$(\forall x_1.F_1 \supset \dots \supset \forall x_n.(F_n \supset A')) \dots$$

then this sequence is equivalent to the rule

$$\frac{\Sigma; \Gamma \longrightarrow F'_1 \quad \dots \quad \Sigma; \Gamma \longrightarrow F'_n}{\Sigma; \Gamma \longrightarrow A} \text{ backchain}$$

which has the proviso that for some Σ -terms t_1, \dots, t_n that have the same types as x_1, \dots, x_n , respectively, it is the case that A is equal to $A'[t_1/x_1, \dots, t_n/x_n]$ and, for $1 \leq i \leq n$, F'_i is equal to $F_i[t_1/x_1, \dots, t_i/x_i]$.

The logic that we have described has been given an efficient implementation in the *Teyjus* system [9]. It is possible also to reason in sophisticated ways about specifications that are constructed using it. To begin with, the logic has strong meta-theoretic properties arising from the fact that derivability in it corresponds exactly to intuitionistic provability. Moreover, it is possible to construct logics incorporating mechanisms such as induction to reason powerfully about what does and does not follow from a given specification [1, 8, 10, 14]. In fact, systems such as Abella [7] and Tac [2] have been constructed to provide computer support for such reasoning.

3. Logic Programming Using the Twelf Specification Language

There are three categories of expressions in LF: *kinds*, *types* or *type families* that are classified by kinds and *objects* or *terms* that are classified by types. We assume two denumerable sets of variables, one for objects and the other for types. We use x and y to denote object variables, u and v to denote type variables and w to denote either. Letting K range over kinds, A and B over types, and M and N over object terms, the syntax of LF expressions is given by the following rules:

$$\begin{aligned} K &:= \text{Type} \mid \Pi x:A.K \\ A &:= u \mid \Pi x:A.B \mid \lambda x:A.B \mid A \ M \\ M &:= x \mid \lambda x:A.M \mid M \ N \end{aligned}$$

Expressions of any of these kinds will be denoted by P and Q . Here, Π and λ are operators that associate a type with a variable

¹Other, non-interpreted type constructors can be added but are not discussed here for simplicity.

²For the reader unfamiliar with such presentations, the expression $\Sigma; \Gamma \xrightarrow{D} A$ corresponds essentially to the selection of the program clause D as the one to backchain on. This then leads to instantiations of universally quantified variables and to the solution of the “body” goals of the clause using the rules $\forall L$ and $\supset L$, culminating eventually in solving the atomic goal by matching it with the head of the clause using the *init* rule.

$$\begin{array}{c}
\frac{}{\Sigma; \Gamma \longrightarrow \top} \top R \quad \frac{\Sigma; \Gamma \cup \{D\} \longrightarrow G}{\Sigma; \Gamma \longrightarrow D \supset G} \supset R \quad \frac{c \notin \Sigma \quad \Sigma \cup \{c\}; \Gamma \longrightarrow G[c/x]}{\Sigma; \Gamma \longrightarrow \forall x. G} \forall R \\
\\
\frac{D \in \Gamma \quad \Sigma; \Gamma \xrightarrow{D} A}{\Sigma; \Gamma \longrightarrow A} \text{decide} \quad \frac{}{\Sigma; \Gamma \xrightarrow{A} A} \text{init} \\
\\
\frac{t \text{ is a } \Sigma\text{-term} \quad \Sigma; \Gamma \xrightarrow{D[t/x]} A}{\Sigma; \Gamma \xrightarrow{\forall x. D} A} \forall L \quad \frac{\Sigma; \Gamma \longrightarrow G \quad \Sigma; \Gamma \xrightarrow{D} A}{\Sigma; \Gamma \xrightarrow{G \supset D} A} \supset L
\end{array}$$

Figure 1. Derivation rules for the *hohh* logic

and bind its free occurrences over the expression after the period. Terms that differ only in the names of bound variables are identified. As with the *hohh* logic, $P[N_1/x_1, \dots, N_n/x_n]$ denotes a simultaneous substitution with renaming to avoid variable capture. We write $A \rightarrow P$ for $\Pi x:A. P$ when x does not appear free in P . We abbreviate $\Pi x_1:A_1. \dots \Pi x_n:A_n. P$ by $\Pi x:\vec{A}. P$.

LF expressions are equipped with a notion of β -reduction defined through the rule $(\lambda x:A. P) N \rightarrow_\beta P[N/x]$. All LF expressions that are well-formed in the sense formalized below normalize strongly under this reduction relation [11]. Moreover any well-typed expression P has a unique normal form up to changes in bound variable names. We denote this normal form by P^β .

The type correctness of LF expressions is assessed relative to contexts that are finite collections of assignments of types and kinds to variables. Formally, contexts, denoted by Γ , are given by the rule

$$\Gamma := \cdot \mid \Gamma, u : K \mid \Gamma, x : A$$

Here, \cdot denotes the empty collection. We write $\text{dom}(\Gamma)$ to denote the variables with assignments in Γ . We are concerned with assertions of the following four forms:

$$\vdash \Gamma \text{ ctx} \quad \Gamma \vdash K \text{ kind} \quad \Gamma \vdash A : K \quad \Gamma \vdash M : A$$

The first assertion signifies that Γ is a well-formed context. The remaining assertions mean respectively that, relative to a (well-formed) context Γ , K is a well-formed kind, A is a well-formed type of kind K and M is a well-formed object of type A . Figure 2 presents the rules for deriving such assertions. Notice that for a context to be well-formed it must not contain multiple assignments to the same variable. To adhere to this requirement, bound variable renaming may be entailed in the use of the *pi-kind*, *pi-fam*, *abs-fam* and *abs-obj* rules. The inference rules allow for the derivation of an assertion of the form $\Gamma \vdash M : A$ only when A is in normal form. To verify such an assertion when A is not in normal form, we first derive $\Gamma \vdash A : \text{Type}$ and then verify $\Gamma \vdash M : A^\beta$. A similar observation applies to $\Gamma \vdash A : K$.

A variable w that appears in an LF expression P that is well-formed with respect to a context Γ has a kind or type of kind *Type* associated with it through either an assignment in Γ or a binding operator. Moreover, the normal form of this kind or type must have a prefix of Π s. If the length of this prefix is n , then an occurrence of w is *fully applied* if it appears in a subterm of the form $w M_1 \dots M_n$. Further, P is *canonical* with respect to Γ if it is in normal form and if every variable occurrence in it is fully applied. A well-formed context Γ is *canonical* if the type or kind it assigns to each variable is canonical relative to Γ . A well-formed type of the form $u M_1 \dots M_n$ that is fully applied is called a *base type*. The LF system admits a notion of η -expansion using which any well-formed expression can be converted into a canonical form.

In later sections we shall consider LF derivations in which all expressions in the end assertion are in normal form. Notice that every expression in the entire derivation must then also be in such a form. This in turn means that in judgments of the forms

$(\lambda x:A. B) : (\Pi x:A'. K)$ and $(\lambda x:A. M) : (\Pi x:A'. B)$ it must be the case that A and A' are identical. Finally, normalization need not be considered in the use of the *var-fam* and *var-obj* rules.

The following “transitivity” property for LF derivations that follows easily from the results in [11] will be useful later; here α stands for any judgment, and substitution and normalization over α and Γ corresponds to distributing these operations to the expressions appearing in them.

Proposition 1 (Substitution). *Let Γ_1, Γ_2 be canonical contexts, and A be a type in canonical form. If $\Gamma_1 \vdash M : A$ has a derivation, and $\Gamma_1, x : A, \Gamma_2 \vdash \alpha$ has a derivation, then $\Gamma_1, (\Gamma_2[M/x])^\beta \vdash (\alpha[M/x])^\beta$ has a derivation as well.*

Additionally we will use a second property of LF derivations, which follows from Proposition 1.

Proposition 2 (Renaming). *Let P be a canonical type or kind, $\Gamma = \Gamma_1, x : P, \Gamma_2$ be a canonical context, and α a canonical judgment. Let y be a variable not bound in Γ , and not occurring in α . Then $\Gamma_1, x : P, \Gamma_2 \vdash \alpha$ has a derivation if and only if $\Gamma_1, y : P, \Gamma_2[y/x] \vdash \alpha[y/x]$ has one.*

The logic programming interpretation of LF is based on viewing types as formulas. More specifically, a specification or program in this setting is given by a context. This starting context, also called a *signature*, essentially describes the vocabulary for constructing types and asserts the existence of particular inhabitants for some of these types. Against this backdrop, questions can be asked about the existence of inhabitants for certain other types. Formally, this amounts to asking if an assertion of the form $\Gamma \vdash M : A$ has a derivation. However, the object M is left unspecified—it is to be extracted from a successful derivation. Thus, the search for a derivation of the assertion is driven by the structure of A and the types available from the context.

A concrete illustration of the paradigm is useful for later discussions.³ Consider a signature or context Γ comprising the following assignments in sequence:

```

nat : Type
z : nat
s : nat → nat
list : Type
nil : list
cons : nat → list → list
append : list → list → list → Type
appNil :  $\Pi K:\text{list}. \text{append nil } K \ K$ 
appCons :  $\Pi X:\text{nat}. \Pi L:\text{list}. \Pi K:\text{list}. \Pi M:\text{list}.$ 

```

³The example of appending lists has been chosen here for its conciseness and because it allows for an easy connection with more traditional forms of logic programming. The primary application domain of Twelf is in specifying (and reasoning about) formal systems such as evaluators and interpreters for languages, type assignment calculi and proof systems. This orientation informs the choice of benchmarks used in Section 6.

$$\begin{array}{c}
\frac{}{\vdash \cdot \text{ ctx}} \text{ null-ctx} \\
\frac{\Gamma \vdash K \text{ kind} \quad \vdash \Gamma \text{ ctx} \quad u \notin \text{ dom}(\Gamma)}{\vdash \Gamma, u : K \text{ ctx}} \text{ kind-ctx} \\
\frac{\Gamma \vdash A : \text{ Type} \quad \vdash \Gamma \text{ ctx} \quad x \notin \text{ dom}(\Gamma)}{\vdash \Gamma, x : A \text{ ctx}} \text{ type-ctx} \\
\frac{\vdash \Gamma \text{ ctx}}{\Gamma \vdash \text{ Type kind}} \text{ type-kind} \quad \frac{\Gamma \vdash A : \text{ Type} \quad \Gamma, x : A \vdash K \text{ kind}}{\Gamma \vdash \Pi x:A. K \text{ kind}} \text{ pi-kind} \\
\frac{\vdash \Gamma \text{ ctx} \quad u : K \in \Gamma}{\Gamma \vdash u : K^\beta} \text{ var-fam} \quad \frac{\vdash \Gamma \text{ ctx} \quad x : A \in \Gamma}{\Gamma \vdash x : A^\beta} \text{ var-obj} \\
\frac{\Gamma \vdash A : \text{ Type} \quad \Gamma, x : A \vdash B : \text{ Type}}{\Gamma \vdash (\Pi x:A. B) : \text{ Type}} \text{ pi-fam} \\
\frac{\Gamma \vdash A : \text{ Type} \quad \Gamma, x : A \vdash B : K}{\Gamma \vdash (\lambda x:A. B) : (\Pi x:A^\beta. K)} \text{ abs-fam} \quad \frac{\Gamma \vdash A : \Pi x:B. K \quad \Gamma \vdash M : B}{\Gamma \vdash (A M) : (K[M/x])^\beta} \text{ app-fam} \\
\frac{\Gamma \vdash A : \text{ Type} \quad \Gamma, x : A \vdash M : B}{\Gamma \vdash (\lambda x:A. M) : (\Pi x:A^\beta. B)} \text{ abs-obj} \quad \frac{\Gamma \vdash M : \Pi x:A. B \quad \Gamma \vdash N : A}{\Gamma \vdash (M N) : (B[N/x])^\beta} \text{ app-obj}
\end{array}$$

Figure 2. Rules for Inferring LF Assertions

$$\begin{array}{l}
(\text{append } L \ K \ M) \rightarrow \\
(\text{append } (\text{cons } X \ L) \ K \ (\text{cons } X \ M))
\end{array}$$

We can ask if there is some term M such that the judgment

$$\begin{array}{l}
\Gamma \vdash M : \text{append } (\text{cons } z \ \text{nil}) \\
\quad (\text{cons } (s \ z) \ \text{nil}) \\
\quad (\text{cons } z \ (\text{cons } (s \ z) \ \text{nil}))
\end{array}$$

is derivable. Assuming that Γ is given by the ambient environment, such a query can be posed in Twelf simply by presenting the type expression. The logic programming interpreter of Twelf will find that the proof term

$$\begin{array}{l}
(\text{appCons } z \ \text{nil} \ (\text{cons } (s \ z) \ \text{nil}) \\
\quad (\text{cons } (s \ z) \ \text{nil}) \\
\quad (\text{appNil } (\text{cons } (s \ z) \ \text{nil})))
\end{array}$$

inhabits this type and hence will succeed on the query. In reaching this conclusion, the interpreter will use the types involving *append* that are present in Γ . Further it will do this in a way that bears a close resemblance to the use of clauses in a Prolog-like setting, interpreting Π like a universal quantifier and \rightarrow like an implication.

The simple example we have considered here will suffice to illustrate most of the later ideas in this paper but it does not bring out the richness of dependent types in specifications. We leave this demonstration to the many discussions already in the literature. We also note that Twelf has many additional features like allowing Π quantification in types to be left implicit and permitting instantiatable variables in queries whose values are to be found through unification. While these aspects are treated in our implementation, to keep the theoretical discussions focused, we shall assume that the only capability that is to be emulated is that of determining the derivability of an assertion of the form $\Gamma \vdash M : A$ in which Γ and A are in canonical form (and M is left unspecified). This assumption is easily justified: these will be “type-checked” prior to conducting a search and the Twelf system assumes equality under η -conversion.

4. From Twelf Specifications to Predicate Formulas

Felty has previously shown how to translate LF specifications and judgments into *hohh* formulas [5, 6]. Her translation proceeds in two steps. First, she describes a coarse mapping of LF expressions into (simply typed) λ -terms. This mapping loses information about dependencies in types and kinds and also does not reflect the correspondences between objects and types and types and kinds. These relationships are encoded later through binary predicates over λ -terms.

The general structure of Felty’s translation is applicable in the context of interest to us. However, the details of her mapping do not quite fit our needs because of her focus on *derivations* in the LF and *hohh* logics. One manifestation of this is that her translation is not based exclusively on types, but assumes also the availability of the objects they are intended to qualify. This is not acceptable in the context of proof search where the task is precisely to determine the existence of those objects: we need a translation that is only based on the type, and which can be applied to an *hohh* metavariable to correspond to an LF query whose object is left unspecified as a metavariable. Second, the correctness result only states an equivalence between LF derivability and *hohh* derivability for *known* LF assertions, and does not consider, for example, whether it is possible for non-canonical or ill-formed objects to be produced in the course of searching for proofs from the *hohh* specification. In contrast, our completeness result will guarantee that after running a query with a metavariable standing for the (encoding of the) object, the only possible instantiations of that metavariable are actual encodings of terms.

The first step towards producing a translation into *hohh* that can be used to interpret Twelf specifications is to adapt Felty’s translation in a way that makes it acceptable in logic programming discussions. Our translation shall only account for judgments of the form $\Gamma \vdash M : A$ since these are the only ones of interest in the logic programming setting described in the previous section. The adequacy of this restriction actually relies on an auxiliary, easily verified, fact: if $\Gamma \vdash A : \text{ Type}$ is known to have a derivation and the last rule in a purported derivation of $\Gamma \vdash M : A$ is an *abs-obj*,

$$\begin{aligned}
\phi(A) &:= \text{lf-obj when } A \text{ is a base type} \\
\phi(\Pi x:A.P) &:= \phi(A) \rightarrow \phi(P) & \phi(\text{Type}) &:= \text{lf-type} \\
\langle u M_1 \dots M_n \rangle &:= u \langle M_1 \rangle \dots \langle M_n \rangle \\
\langle x M_1 \dots M_n \rangle &:= x \langle M_1 \rangle \dots \langle M_n \rangle & \langle \lambda x:A.M \rangle &:= \lambda^{\phi(A)} x. \langle M \rangle \\
\llbracket \Pi x:A.B \rrbracket &:= \lambda M. \forall x. (\llbracket A \rrbracket x) \supset (\llbracket B \rrbracket (M x)) \\
\llbracket A \rrbracket &:= \lambda M. \text{hastype } M \langle A \rangle \text{ where } A \text{ is a base type}
\end{aligned}$$

Figure 3. Encoding of types, objects, and simplified translation of LF judgments to *hohh*

$$\begin{aligned}
&\text{hastype } z \text{ nat} \\
&\forall n. \text{hastype } n \text{ nat} \supset \text{hastype } (s \ n) \text{ nat} \\
&\text{hastype nil list} \\
&\forall n. \text{hastype } n \text{ nat} \supset \forall l. \text{hastype } l \text{ list} \supset \text{hastype } (\text{cons } n \ l) \text{ list} \\
&\forall l. \text{hastype } l \text{ list} \supset \text{hastype } (\text{appNil } l) (\text{append nil } l) \\
&\forall x. \text{hastype } x \text{ nat} \supset \forall l. \text{hastype } l \text{ list} \supset \forall k. \text{hastype } k \text{ list} \supset \forall m. \text{hastype } m \text{ list} \supset \\
&\quad \forall a. \text{hastype } a (\text{append } l \ k \ m) \supset \text{hastype } (\text{appCons } x \ l \ k \ m \ a) (\text{append } (\text{cons } x \ l) \ k (\text{cons } x \ m))
\end{aligned}$$

Figure 4. Simple translation of the LF specification for *append*

then the left premise for the latter derivation must have a proof and hence does not need to be encoded by the translation.

Our translation is presented in Figure 3. This translation first encodes LF objects and types in *hohh* terms by dropping a lot of typing information; as mentioned already, this information will be recovered later in the encoding of LF judgments. Under this translation, an object (type) of type (kind) P is represented by an *hohh* term of simple type $\phi(P)$, built from the atomic types *lf-type* and *lf-obj*. The encoding of an object or base type Q is then given by $\langle Q \rangle$; note that in the process we assume a reuse of (LF) variable names with an appropriate type as part of the corresponding *hohh* signature. As an example, the LF signature at the end of the last section leads to the following *hohh* signature:

$$\begin{aligned}
&\text{nat} : \text{lf-type} \\
&z : \text{lf-obj} \\
&s : \text{lf-obj} \rightarrow \text{lf-obj} \\
&\text{list} : \text{lf-type} \\
&\text{nil} : \text{lf-obj} \\
&\text{cons} : \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \text{lf-obj} \\
&\text{append} : \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \text{lf-type} \\
&\text{appNil} : \text{lf-obj} \rightarrow \text{lf-obj} \\
&\text{appCons} : \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \\
&\quad \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \text{lf-obj} \rightarrow \text{lf-obj}
\end{aligned}$$

Further, the LF type *append nil nil nil* gets translated to the same term in *hohh*, where it has type *lf-type*. This translation behaves well with respect to substitution and β -conversion, and is injective for objects (types) of the same type (kind). Finally, we take up the translation of LF type assignments and judgments in the last two clauses in Figure 3. To emphasize reliance only on the structure of types, these clauses describe explicitly only the translation of an LF type A . Such a type is mapped into an *hohh* predicate denoted by $\llbracket A \rrbracket$ that, intuitively, codifies the property of being a translation of an LF object of type A . This translation is defined on all canonical types and uses the *hohh* predicate *hastype* of type *lf-obj* \rightarrow *lf-type* \rightarrow *o*. If A is a base type, $\llbracket \Pi x_1:B_1 \dots \Pi x_n:B_n.A \rrbracket$ has type $\tau \rightarrow o$ where τ is *lf-obj* $\rightarrow \dots \rightarrow \text{lf-obj} \rightarrow \text{lf-obj}$ with n negative occurrences of *lf-obj*. Once the translation of LF types is in place, we define $\llbracket M : A \rrbracket$ derivatively to be $(\llbracket A \rrbracket \langle M \rangle)$.

Twelf specifications are encoded by dropping all kind assignments and translating each type assignment they contain. As an ex-

ample, the Twelf specification of *append* translates into the clauses in Figure 4. From these clauses, we can, for example, derive the goal *hastype* (*cons* (*s z*) *nil*) *list* and we could search for terms X satisfying the goal

$$\begin{aligned}
&\text{hastype } X (\text{append } (\text{cons } z \ \text{nil}) \\
&\quad (\text{cons } (s \ z) \ \text{nil}) \\
&\quad (\text{cons } z \ (\text{cons } (s \ z) \ \text{nil}))).
\end{aligned}$$

Let Γ' be the translation of an LF context Γ and α' be the translation of the LF judgment α . These translations are based on an implicit *hohh* signature Σ . In the case that all the free variables in α belong to $\text{dom}(\Gamma)$, then, in fact, Σ consists of an isomorphic copy of the symbols in $\text{dom}(\Gamma)$. Henceforth, we shall assume Σ to be just such an *hohh* signature and we shall write $\Gamma' \rightarrow \alpha'$ as a shorthand for $\Sigma; \Gamma' \rightarrow \alpha'$. The correctness of the (simple) translation is then the content of the following theorem.

Theorem 1. *Let Γ be a well-formed canonical LF context and let A be a canonical LF type such that $\Gamma \vdash A : \text{Type}$ has a derivation. If $\Gamma \vdash M : A$ has a derivation for a canonical object M , then there is a derivation of $\llbracket \Gamma \rrbracket \rightarrow \llbracket M : A \rrbracket$. Conversely, if $\llbracket \Gamma \rrbracket \rightarrow (\llbracket A \rrbracket \ M)$ has a derivation for any *hohh* term M of appropriate type, then there is a canonical LF object M' such that $M = \langle M' \rangle$ and $\Gamma \vdash M' : A$ has a derivation.*

Proof outline Completeness can be proved by a simple induction on the LF derivation, building an *hohh* derivation that mimics its structure. Soundness is more involved: we proceed by induction on the *hohh* derivation, gradually recovering the structure of M' , maintaining the derivability of $\Gamma \vdash A : \text{Type}$ that allows us to build an LF derivation even in the case that *abs-obj* was the last rule used. The detailed proof is presented in Appendix A.

The simple translation presented in this section cannot be the basis of a practical implementation of logic programming in LF. Proof search using a program it produces may involve repeatedly proving goals of the form *hastype* $M \ A$ for (encodings of) the same object M and type A . This can be seen from the example in Figure 4: at every step of deriving an instance of *append*, the lists must be checked to be well-typed, which artificially introduces a quadratic complexity. An important point to note, however, is that this redundancy in “type-checking” is not easily detectable from the *hohh* program that is generated. Rather, it must be determined,

$$\begin{array}{c}
\frac{\Gamma; \cdot; x \sqsubseteq_o A_i \text{ for some } A_i}{\Gamma; x \sqsubseteq_t c\vec{A}} \text{APP}_i \\
\\
\frac{y_i \in \delta \text{ for each } y_i \quad y_i \text{ distinct}}{\Gamma; \delta; x \sqsubseteq_o x \vec{y}} \text{INIT}_o \\
\\
\frac{\Gamma, y; x \sqsubseteq_t B}{\Gamma; x \sqsubseteq_t \Pi y:A.B} \text{PI}_i \\
\\
\frac{y \notin \Gamma \text{ and } \Gamma; \delta; x \sqsubseteq_o M_i \text{ for some } i}{\Gamma; \delta; x \sqsubseteq_o y \vec{M}} \text{APP}_o \\
\\
\frac{\Gamma; \delta, y; x \sqsubseteq_o M}{\Gamma; \delta; x \sqsubseteq_o \lambda y:A.M} \text{ABS}_o
\end{array}$$

Figure 5. Rigidly occurring variables in types and objects

and shown to be safely eliminable, based on deeper properties of LF terms. It is this issue that we take up in the next section.

5. An Improved Translation of Twelf Specifications

In order to make the translation of LF specifications into *hohh* practical from an implementation standpoint, we make two optimizations.

The first, and main, optimization exploits the fact that we are considering derivations of the form $\Gamma \vdash M : A$ where Γ and A have already been type-checked. For example, we may be wanting to determine whether the LF type

$$\text{append}(\text{cons } z \text{ nil}) \text{ nil} (\text{cons } z \text{ nil})$$

is inhabited. Before attempting to do this, we would have already determined that $\text{append}(\text{cons } z \text{ nil}) \text{ nil} (\text{cons } z \text{ nil})$ is a valid type, which means, for instance, that we would have checked that $(\text{cons } z \text{ nil})$ is a valid object of type *list*. Therefore, there is no need to show again that $(\text{cons } z \text{ nil})$ has this property in the course of searching for an inhabitant of the displayed type. Our optimized translation takes advantage of this kind of observation by statically removing some run-time checking from the translation of LF typing. More specifically, our optimization is based on the following idea. Suppose we can determine that, for a particular i , t_i must always appear in the type $(A[t_1/x_1, \dots, t_n/x_n])^\beta$. Then the translation of the type $\Pi x_1:B_1 \dots \Pi x_n:B_n.A$ does not need to include explicit type-checking over the instantiation of x_i . We characterize some of these cases by using the notion of a *rigid occurrence* of x_i in A that is expressed formally through the judgment $\vec{x}; x_i \sqsubseteq_t A$ defined by the rules in Figure 5; the rules APP_i and PI_i in this figure act on LF types, and the rules INIT_o , APP_o , and ABS_o act on LF objects. We shall allow type checking over instantiations of rigid variables to be eliminated from the simple translation. By doing so, we shall both reap an efficiency benefit and also make the result of translation correspond more closely to the original LF type.

The second optimization is more transparent, not depending on deep properties of dependent types. The essential observation is the following. Instead of producing predicates of the form

$$\text{hastype } X (\text{append } L \text{ } K \text{ } M)$$

and *hastype* L *list*, we can specialize them to $\text{append } X \text{ } L \text{ } K \text{ } M$ and *list* L . This results in a *hohh* program that is much clearer, and more closely related to the original LF specification. Moreover, this simple transformation can also lead to better performance in a logic programming setting because it allows for the exploitation of a common optimization, namely, the indexing on a predicate name

that speeds up the determination of candidate clauses on which to backtrack.

The improved translation that uses these two ideas is presented on Figure 6. The $\llbracket \bullet \rrbracket_\Gamma^+$ translation is used on type assignments appearing negatively (notably context items) and $\llbracket \bullet \rrbracket^-$ on positive typing judgments (notably the conclusion of LF assertions). As before, that translation is entirely guided by the type, and defined for all canonical types. We shall use the notation $\llbracket M : A \rrbracket^-$ for $(\llbracket A \rrbracket^-(M))$, and define $\llbracket \Gamma \rrbracket^+$ as the result of applying $\llbracket \bullet \rrbracket_\Gamma^+$ to each context item, dropping kind assignments. Note that instead of replacing unnecessary typing judgments with \top we could simply elide them all together; we use \top as a placeholder because it simplifies later proofs. This translation is illustrated by its application to the example Twelf specification considered in Section 3 that yields the clauses shown in Figure 7. These clauses should be contrasted with the ones in Figure 4 that are produced by the earlier, naive translation.

We shall now establish the correctness of the optimized translation. We first prove a fundamental lemma concerning rigidly occurring variables, that is in fact an observation about LF: for an LF base type A , if we have derivations of

$$\begin{array}{l}
\Gamma \vdash \Pi x_1:B_1 \dots \Pi x_n:B_n.A : \text{Type} \quad \text{and} \\
\Gamma \vdash A[t_1/x_1 \dots t_n/x_n] : \text{Type}
\end{array}$$

and there is a rigid occurrence of x_i in A , i.e., $\vec{x}; x_i \sqsubseteq_t A$ has a derivation, then $\Gamma \vdash t_i : B_i[t_1/x_1 \dots t_{i-1}/x_{i-1}]$ has a derivation. The idea of the proof is as follows. The judgment $\vec{x}; x_i \sqsubseteq_t A$ gives a path in A that leads to x_i , and this path can never be erased by the considered substitution; following this path simultaneously in the two LF derivations, one eventually finds on one side a derivation of $\Gamma \vdash x_i : B_i$ and on the other side the expected derivation of $\Gamma \vdash t_i : B_i[t_1/x_1, \dots, t_{i-1}/x_{i-1}]$.

In order to be able to use this observation in our correctness argument, we formulate a stronger, rather technical lemma that deals directly with encoded types that are the result of instantiations of (a priori) arbitrary *hohh* terms, and ensures that discovered *hohh* terms are in fact encodings of LF objects. These technical details concerning encodings are tedious but shallow, and the essential structure of the proof follows the lines sketched above.

Definition 1. Let \vec{t} be a vector of *hohh* terms, and \vec{x} a vector of variables of the same length. If M and N are LF objects, then we write $(M \sim N)[t_1/x_1 \dots t_n/x_n]$ when

$$\langle M \rangle = \langle N \rangle[t_1/x_1 \dots t_n/x_n].$$

For LF types A and B , we write $(A \sim B)[t_1/x_1 \dots t_n/x_n]$ when the two types are equal up to $(\bullet \sim \bullet)[t_1/x_1 \dots t_n/x_n]$ on objects within. Finally we extend this notion to contexts of the same length by pushing it down to the types bound by the context. We shall omit \vec{t} and \vec{x} when they are obvious from the context, simply writing $P \sim Q$.

Lemma 1. Let \vec{t} be a vector of *hohh* terms, \vec{x} a vector of variables, and \vec{B} of canonical LF types, all of same length, such that $t_j = \langle t'_j \rangle$ for $j < i$. Let $\Gamma_0 = x_1 : B_1, \dots, x_n : B_n$.

1. Let Γ and Δ be LF contexts, M an LF object and A a type, all being assumed canonical. Let δ be $\text{dom}(\Delta)$. Suppose that there are derivations of $\vec{x}; \delta; x_i \sqsubseteq_o M$ and $\Gamma, \Gamma_0, \Delta \vdash M : A$ and $\Gamma, \Delta' \vdash M' : A'$, with $A' \sim A$, $M' \sim M$ and $\Delta' \sim \Delta$. Then t_i is of the form $\langle t'_i \rangle$ and there is a derivation of $\Gamma \vdash t'_i : B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}]$.
2. Let $\Pi x:\vec{B}.A$ be a canonical type, where A is a base type. Suppose that $\Gamma \vdash \Pi x:\vec{B}.A : \text{Type}$ and $\vec{x}; x_i \sqsubseteq_t A$ have derivations. Further, for some A' such that $A' \sim A$, suppose

$$\begin{aligned}
\llbracket \Pi x:A.B \rrbracket_{\Gamma}^+ &:= \begin{cases} \lambda M. \forall x. \top \supset \llbracket B \rrbracket_{\Gamma, x}^+(M x) & \text{if } \Gamma; x \sqsubset_t B \\ \lambda M. \forall x. \llbracket A \rrbracket_{\Gamma}^-(x) \supset \llbracket B \rrbracket_{\Gamma, x}^+(M x) & \text{otherwise} \end{cases} \\
\llbracket u \vec{N} \rrbracket_{\Gamma}^+ &:= \lambda M. u M \langle \vec{N} \rangle \\
\llbracket \Pi x:A.B \rrbracket_{\Gamma}^- &:= \lambda M. \forall x. \llbracket A \rrbracket_{\Gamma}^+(x) \supset \llbracket B \rrbracket_{\Gamma, x}^-(M x) \\
\llbracket u \vec{N} \rrbracket_{\Gamma}^- &:= \lambda M. u M \langle \vec{N} \rangle
\end{aligned}$$

Figure 6. Optimized translation of LF specifications and judgments to *hohh*

$$\begin{aligned}
& \text{nat } z \\
& \forall n. \text{nat } n \supset \text{nat } (s n) \\
& \text{list nil} \\
& \forall n. \text{nat } n \supset \forall l. \text{list } l \supset \text{list } (\text{cons } n l) \\
& \forall l. \top \supset \text{append } (\text{appNil } l) \text{ nil } l l \\
& \forall x. \top \supset \forall l. \top \supset \forall k. \top \supset \forall m. \top \supset \\
& \quad \forall a. \text{append } a l k m \supset \text{append } (\text{appCons } x l k m a) (\text{cons } x l) k (\text{cons } x m)
\end{aligned}$$

Figure 7. Optimized translation of the LF specification for *append*

that $\Gamma \vdash A' : \text{Type}$ has a derivation. Then $t_i = \langle t'_i \rangle$ and there is a derivation of $\Gamma \vdash t'_i : B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}]$.

Proof. We prove part (1) by induction on the structure of the derivation of $\vec{x}; \delta; x_i \sqsubset_o M$. In the argument below, we let \mathcal{D} be the derivation of $\Gamma, \Gamma_0, \Delta \vdash M : A$, and \mathcal{D}' be the derivation of $\Gamma, \Delta' \vdash M' : A'$.

- In the base case of INIT_o , $M = x_i \vec{y}$ where \vec{y} are distinct bound variables from δ . The derivation \mathcal{D} must consist of n *app-obj* rules and a *var-obj* rule on x_i , whose type B_i must be of the form $\Pi z:\vec{C}.D$, with $A = D[\vec{y}/\vec{z}]$. Note that, because the variables y_i are distinct bound variables that are fresh with respect to D , this substitution can be inverted, and we thus have $A[\vec{z}/\vec{y}] = D$. The other subderivations of the chain of *app-obj* applications are instances of *var-obj* establishing $y_i : C_i[\vec{y}/\vec{z}]$, hence $(y_i : C'_i[\vec{y}/\vec{z}]) \in \Delta'$ for $C'_i \sim C_i$. We next determine t'_i . By η -equivalence we can assume that t_i is of the form $\lambda z_1 \dots \lambda z_n. u$. We have

$$\langle M' \rangle = t_i \vec{y} = u[\vec{y}/\vec{z}],$$

hence $u = \langle M' \rangle[\vec{z}/\vec{y}] = \langle M'[\vec{z}/\vec{y}] \rangle$. Let $u' = M'[\vec{z}/\vec{y}]$ and $t'_i = \lambda z:\vec{C}'.u'$. We have

$$\begin{aligned}
\langle t'_i \rangle &= \lambda z_1 \dots \lambda z_n. \langle M' \rangle[\vec{z}/\vec{y}] \\
&= \lambda z_1 \dots \lambda z_n. u = t_i.
\end{aligned}$$

We know that \mathcal{D}' derives $\Gamma, \Delta' \vdash M' : A'$. From this we obtain a derivation of

$$\Gamma, \Delta'[\vec{z}/\vec{y}] \vdash u' : A'[\vec{z}/\vec{y}]$$

by renaming variables \vec{y} into \vec{z} , employing Proposition 2. The context $\Delta'[\vec{z}/\vec{y}]$ contains assignments $(z_i : C'_i)$ and the other variables in its domain do not occur in u' nor $A'[\vec{z}/\vec{y}]$ (since $A' \sim A$, $A = D[\vec{y}/\vec{z}]$ and D is a subterm of B_i which cannot contain any y_i). We then have

$$\Gamma \vdash (\lambda z:\vec{C}'.u') : (\Pi z:\vec{C}'.A'[\vec{z}/\vec{y}])$$

by weakening unused variables and using *abs-obj* to introduce the variables \vec{z} . This is a typing derivation for t'_i ; we must now show that the associated type is actually the expected one:

$$B_i[t'_1/x_1 \dots t'_{i-1}/x_{i-1}]$$

We have $\langle A \rangle[t_1/x_1 \dots t_n/x_n] = \langle A[t'_1/x_1 \dots t'_{i-1}/x_{i-1}] \rangle$ and $A' \sim A$, from which we obtain, by injectivity of $\langle \bullet \rangle$, that $A' = A[t'_1/x_1 \dots t'_{i-1}/x_{i-1}]$. The same goes for C'_i and C_i . Since $B_i = \Pi z:\vec{C}.A[\vec{z}/\vec{y}]$, and the substitutions $[t'_1/x_1 \dots t'_{i-1}/x_{i-1}]$ and $[\vec{z}/\vec{y}]$ permute, we have:

$$\Pi z:\vec{C}'.A'[\vec{z}/\vec{y}] = B_i[t'_1/x_1 \dots t'_{i-1}/x_{i-1}]$$

- In the ABS_o case, we have $M = \lambda y:A_1.N$ and \mathcal{D} ends with the *abs-obj* rule as follows:

$$\frac{\Gamma, \Gamma_0, \Delta \vdash A_1 : \text{Type} \quad \Gamma, \Gamma_0, \Delta, y : A_1 \vdash N : A_2}{\Gamma, \Gamma_0, \Delta \vdash (\lambda y:A_1.N) : (\Pi y:A_1.A_2)}$$

Then $A' \sim \Pi y:A_1.A_2$, and hence A' must be of the form $\Pi y:A'_1.A'_2$ where $A'_i \sim A_i$. Similarly, we obtain that M' is of the form $\lambda y:A'_1.N'$ with $N' \sim N$. Then, \mathcal{D}' must contain a derivation of

$$\Gamma, \Delta', y : A'_1 \vdash N' : A'_2,$$

and we conclude by the inductive hypothesis.

- In the APP_o case, we have $M = y N_1 \dots N_m$, $y \notin \vec{x}$ and $\vec{x}; \delta; x_i \sqsubset_o N_j$. Let $\Pi z_1:C_1 \dots \Pi z_m:C_m.D$ be the type of y in (Γ, Δ) . The derivation \mathcal{D} starts with a chain of *app-obj* applications, followed by *var-obj* on y . The premise corresponding to N_j establishes that

$$\Gamma, \Gamma_0, \Delta \vdash N_j : C_j[N_1/z_1, \dots, N_{j-1}/z_{j-1}]$$

In (Γ, Δ') , the variable y is assigned the type $\Pi z:\vec{C}'.D'$ with all $C'_k \sim C_k$. Moreover, since $M' \sim (y N_1 \dots N_m)$ and since y is not affected by the instantiation of \vec{x} , it must be that M' is of the form $(y N'_1 \dots N'_m)$ with all $N'_j \sim N_j$. The derivation \mathcal{D}' must proceed in a similar fashion, namely a chain of *app-obj* applications followed by *var-obj* on y . Therefore we have a derivation of

$$\Gamma, \Delta' \vdash N'_j : C'_j[N'_1/z_1, \dots, N'_{j-1}/z_{j-1}]$$

We can conclude by the inductive hypothesis because

$$C'_j[N'_1/z_1 \dots N'_{j-1}/z_{j-1}] \sim C_j[N_1/z_1 \dots N_{j-1}/z_{j-1}]$$

(which relies on the disjointness of \vec{x} and \vec{z}).

The proof of (2) follows a similar pattern. First, by a straightforward inspection of the first rules of the derivation of

$$\Gamma \vdash \Pi x:\vec{B}.A : \text{Type}$$

we extract a derivation of $\Gamma, \Gamma_0 \vdash A : \text{Type}$. Then, since A is a base type, it must be (by rule APP_i) that x_i rigidly occurs in one of its arguments M . Note that A and A' have the same structure on the path leading to M , since no object is involved there. Hence, a simultaneous inspection of the first rules of the derivations of $\Gamma, \Gamma_0 \vdash A : \text{Type}$ and $\Gamma \vdash A' : \text{Type}$ yields derivations of $\Gamma, \Gamma_0 \vdash M : T$ and $\Gamma \vdash M' : T'$ for $M' \sim M$ and $T' \sim T$. We can conclude using part (1). \square

The definition of rigidity described above might seem restrictive. In particular, one might want to allow

$$\Gamma; \delta; x \sqsubset_o x \vec{N}$$

in INIT_o . However, with such a rule the rigidity lemma described above is no longer true. For example, in a signature Γ containing $\text{num} : \text{nat} \rightarrow \text{Type}$ and $\text{num}_n : \Pi n:\text{nat}.(\text{num } n)$, the object $t = \text{num}_n$ provides a counter-example to Lemma 1, part (1): we have $\Gamma, x : (\text{nat} \rightarrow \text{num } z) \vdash (x z) : (\text{num } z)$ and $\Gamma \vdash (t z) : (\text{num } z)$ but not $\Gamma \vdash t : \text{nat} \rightarrow \text{num } z$. This example highlights a crucial aspect of our definition: the applications allowed in INIT_o should always induce *invertible* substitutions. As in higher-order pattern unification [13, 18], we achieve this by restricting to applications involving a simple form of β -reductions that are similar to renaming.

We now use Lemma 1 to prove the correctness of the optimized translation.

Theorem 2. *Let Γ be an LF context, A an LF type, both canonical, such that $\vdash \Gamma \text{ ctx}$ and $\Gamma \vdash A : \text{Type}$ are derivable. Then when M is an arbitrary hohh term, $\llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket(M)$ has a derivation if and only if $\llbracket \Gamma \rrbracket^+ \rightarrow \llbracket A \rrbracket^-(M)$ has a derivation.*

Proof. We establish the soundness direction by induction on the derivation of the optimized translation, maintaining the assumptions about Γ and A .

If A is of the form $\Pi x:B.A'$ our derivation ends as follows:

$$\frac{\llbracket \Gamma, x : B \rrbracket^+ \rightarrow \llbracket A' \rrbracket^-(M x)}{\llbracket \Gamma \rrbracket^+ \rightarrow \llbracket \Pi x:B.A' \rrbracket^-(M)} \forall R, \supset R$$

First, $\Gamma \vdash B : \text{Type}$, $\vdash (\Gamma, x : B) \text{ ctx}$ and $\Gamma, x : B \vdash A' : \text{Type}$ must have derivations since Γ and A are well-formed. We can thus apply the inductive hypothesis, obtaining that

$$\llbracket \Gamma, x : B \rrbracket \rightarrow \llbracket A' \rrbracket(M x)$$

has a derivation. By $\forall R$ and $\supset R$, $\llbracket \Gamma \rrbracket \rightarrow \llbracket \Pi x:B.A' \rrbracket(M)$ has one as well.

If A is a base type, then our derivation starts with a backchaining on the encoding of some $(y : \Pi x:\vec{B}.A') \in \Gamma$, i.e., on

$$\forall x_1. (\llbracket B_1 \rrbracket^-(x_1) \supset \dots \supset \forall x_n. (\llbracket B_n \rrbracket^-(x_n) \supset (u(y \vec{x}) \overrightarrow{\langle N \rangle}))).$$

In particular, this rule application has the form

$$\frac{\llbracket \Gamma \rrbracket^+ \rightarrow F_1 \quad \dots \quad \llbracket \Gamma \rrbracket^+ \rightarrow F_n}{\llbracket \Gamma \rrbracket^+ \rightarrow (u(y \vec{x}) \overrightarrow{\langle N \rangle})[t/x]} \text{backchain}$$

where F_i is either $(\llbracket B_i \rrbracket^-(x_i))[t_1/x_1, \dots, t_i/x_i]$ or \top . We perform an inner induction on $i \leq n$, showing that for all $j \leq i$, $t_j = \langle t'_j \rangle$ for some LF object t'_j , and that we have derivations of

$$\llbracket \Gamma \rrbracket \rightarrow (\llbracket B_j[t'_1/x_1, \dots, t'_{j-1}/x_{j-1}] \rrbracket t'_j)$$

and

$$\Gamma \vdash t'_j : B_j[t'_1/x_1, \dots, t'_{j-1}/x_{j-1}].$$

- We first treat the case where $F_i = \top$, i.e., there is a derivation of $\vec{x}; x_i \sqsubset_t A'$. We assumed that $\Gamma \vdash A : \text{Type}$, and since Γ is valid we also have a derivation of $\Gamma \vdash \Pi x:\vec{B}.A' : \text{Type}$. We can thus apply Lemma 1, to obtain t'_i and a derivation of $\Gamma \vdash t'_i : B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}]$, and we conclude by Theorem 1.
- When $F_i \neq \top$, we can see that that within the derivation of

$$\Gamma \vdash \Pi x:\vec{B}.A' : \text{Type}$$

there is a derivation of

$$\Gamma, x_1 : B_1, \dots, x_{i-1} : B_{i-1} \vdash B_i : \text{Type}.$$

By substituting (Proposition 1) the derivations provided by the inner inductive hypothesis on this formula we construct a derivation of

$$\Gamma \vdash B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}] : \text{Type}.$$

We can now apply the outer inductive hypothesis on F_i , to conclude that $\llbracket \Gamma \rrbracket \rightarrow (\llbracket B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}] \rrbracket t_i)$ has a derivation. By Theorem 1, we finally obtain that t_i is of the form $\langle t'_i \rangle$.

We compose all derivations

$$\llbracket \Gamma \rrbracket \rightarrow \llbracket B_i[t'_1/x_1, \dots, t'_{i-1}/x_{i-1}] \rrbracket t_i$$

by *backchain* on the encoding of $(y : \Pi x:\vec{B}.A') \in \Gamma$, obtaining the expected derivation of

$$\llbracket \Gamma \rrbracket \rightarrow \text{hastype } (y \vec{t}) (u \overrightarrow{\langle N \rangle})[t/x]$$

Completeness is proved by an induction on the derivation of the simple translation. This direction is rather straightforward as it consists only of dropping information. Details can be found in Appendix A. \square

Therefore, by Theorems 1 and 2, intuitionistic provability under the optimized translation is equivalent to provability in LF, and the following is a theorem.

Theorem 3 (Optimized translation correctness). *Let Γ be an LF specification such that $\vdash \Gamma \text{ ctx}$ has a derivation, A an LF type such that $\Gamma \vdash A : \text{Type}$ has a derivation. Then, for any LF object M such that $\Gamma \vdash M : A$ has a derivation, $\llbracket \Gamma \rrbracket^+ \rightarrow \llbracket M : A \rrbracket^-$ is derivable. Moreover, if $\llbracket \Gamma \rrbracket^+ \rightarrow \llbracket A \rrbracket^-(M)$ for an arbitrary hohh term M , then it must be that $M = \langle M' \rangle$ for some canonical LF object such that $\Gamma \vdash M' : A$ has a derivation.*

6. Performance Comparisons

We have claimed two properties for our translation: that it produces an *hohh* program which corresponds closely to the original LF specification, and that this program provides an effective means for executing the specification. Evidence for the first claim is provided by the translation of the *append* specification presented in Figure 7, especially when one uses the easily applied simplification of a formula of the form $\top \supset F$ to F . Notice also the correspondence of the definition of the *append* predicate to the one that one might in, e.g., Prolog, if one drops the first “proof term” argument of the predicate. To fully appreciate this benefit, it is necessary to consider larger examples that space does not allow us to do in this paper. However, such examples are available with the implementation [23]. We suggest that the reader look especially at the example of the evaluator for Mini-ML with terms that are not indexed by their type that is described below in the collection of benchmarks: the

Example	Twelf	Simple	Optimized	Typed Optimized	Indexing
reverse(10)	1.0	0.40	0.14	0.07	0.08
reverse(20)	1.0	0.57	0.19	0.12	0.11
reverse(30)	1.0	0.63	0.20	0.14	0.11
reverse(40)	1.0	0.41	0.13	0.10	0.07
reverse(50)	1.0	0.46	0.15	0.10	0.08
miniml(50)	1.0	0.74	0.25	0.18	0.08
miniml(100)	1.0	1.25	0.44	0.30	0.17
miniml(150)	1.0	1.75	0.56	0.41	0.25
miniml(200)	1.0	2.89	0.83	0.62	0.41
typed miniml(50)	1.0	2.27	1.07	0.57	0.48
typed miniml(100)	1.0	2.22	0.76	0.49	0.38
typed miniml(150)	1.0	3.49	1.44	0.67	0.55
typed miniml(200)	1.0	3.70	0.92	0.67	0.55
perm(10)	1.0	overflow	3.13	0.94	0.72
perm(20)	1.0	overflow	1.75	0.78	0.44
perm(30)	1.0	overflow	3.05	1.52	0.81
perm(40)	1.0	overflow	3.95	2.15	1.14
perm(50)	1.0	overflow	5.05	2.88	1.59
num(64)	1.0	158.19	0.25	0.23	0.21
num(128)	1.0	∞	0.10	0.10	0.07
num(256)	1.0	∞	0.15	0.14	0.13
num(512)	1.0	∞	0.003	0.003	0.003

Figure 8. Performance comparison results

translation results in an *hohh* program that is what one might write in *hohh* directly.

To test the second claim, we have carried out performance comparisons between the Twelf implementation that interprets LF specifications directly via a Standard ML program and an implementation obtained by translating these specifications into *hohh* programs and then executing these using the Teyjus system. We present results here over programs that have a few different characteristics:

- First, as we are interested in logic programming in LF, the traditional logic program for naively reversing a list a n times is included.
- The encoding of evaluators for various languages is a common usage of LF. We have therefore used an encoding of Mini-ML along with an encoding of addition as another sample program. This benchmark, called *miniml*, consists of adding n to 10 using the encoding.
- The *miniml* specification does not make essential use of dependent types. The *typed miniml* benchmark, which consists of an evaluator for Mini-ML in which terms are indexed by their type, uses dependent types to ensure that terms are well-formed. The Mini-ML program that was run is a typed version of the encoding of addition.
- An implementation of a meta-interpreter for intuitionistic non-commutative linear logic (INCLL) has been proposed as a test program [21]. The *perm* benchmark tests list permutation encoded in INCLL and run using the meta-interpreter on lists of length n .
- The last benchmark, referred to as *num*, involves rewriting arithmetic expressions into an equivalent normal form. This example again makes essential use of dependent types by associating with each equivalence of two such terms a proof of their equivalence. The benchmark tests rewriting expressions of size n .

The third through fifth columns of Figure 8 present data comparing the simple translation, the translation with redundant typing judgments removed, and the fully optimized translation against the standard of Twelf with default optimizations on these bench-

marks.⁴ As described in Figure 6, the fully optimized translation inserts the proof term as the first argument of the predicate generated. Since this term is to be determined by proof search, advantage cannot be taken of the capability Teyjus possesses of indexing on the first argument. The last column presents data for the case where we make the proof term the last argument instead. In the data presented, overflow indicates a heap overflow in the Teyjus simulator, and ∞ means that the program ran more than 1000 times longer than Twelf.

The most optimized translation leads to better performance in most cases, often significantly so. On the other hand, the simple translation yields a program that is generally slower than Twelf. In particular, performance tends to deteriorate with larger problems sizes, in keeping with the difficulty that we noted with this translation. However, the simple translation is still comparable to Twelf on the first three benchmarks. On the *perm* benchmark, Twelf does quite well and even out-performs Teyjus with the optimized translation on problems of large size. We have yet to pinpoint the reason for this—the program is large and difficult to analyze in detail—but we suspect that the linear head optimization that delays expensive unification computation till after simpler checks have been made may have something to do with this. The fact that term indexing causes significant improvement with Teyjus gives credence to this observation.

For problems of very large size with all the benchmarks, the performance of Twelf deteriorates quite dramatically; this is seen, for example, in the case of *num(n)* for a problem of size 512. This phenomenon is linked to the fact that Twelf consumes excessive amounts of memory. The ultimate source of this problem is perhaps the fact that Twelf is implemented in SML: it has been argued that realizing a logic programming language in a functional programming setting can lead to poor memory reclamation and eventually to shortage of space [3].

⁴This setting with Twelf leads to the best performance on these examples.

7. Conclusion and Future Work

We have considered in this paper a translation of Twelf specifications into logic programs in the *hohh* language. An important part of our ideas is the recognition of certain situations in which type information is redundant in LF expressions and hence its checking can be avoided. Our eventual translation produces a program that corresponds closely to the original specification and we have argued that it can be the basis for an effective animation of Twelf descriptions.

The specific work undertaken here can be extended in a few different ways. As an extension to our notion of rigidity, we might observe that, when applying a variable of type $\Pi x:\vec{B}.A$, we could identify redundant type information, not only between a B_i and A , but also between a B_i and a different B_j . It would also be interesting to relate our work to the ideas of Reed [22] who describes a notion of *strictness* similar to rigidity, used for the different purpose of identifying sub-terms of LF objects that could be reconstructed if elided – in contrast, we avoid redundant type checking but still generate a complete LF object. Such an understanding might lead both to an improvement of our translation and to the ability to shorten LF terms that are needed in applications such as that of proof-carrying-code [17]. From an implementation perspective, another possible optimization is to avoid constructing an LF object explicitly when the task has been identified as that of only determining whether a type has an inhabitant: experiments in this direction indicate in some cases a ten-fold performance improvement over the optimized translation. Techniques from the area of extracting programs from proofs that pertain to isolating parts of a proof that do not contribute to its overall computational content—e.g., see [25]—are potentially useful to the application of such an optimization; these techniques might provide the basis for noting components of a type whose inhabitants do not participate in the term corresponding to the overall type.

We have focused here on realizing Twelf through a translation to λ Prolog. A different approach, worthy of investigation, is that of compiling Twelf specifications directly to bytecode for the virtual machine underlying the Teyjus system. Such an approach would make it possible to realize optimizations that have been developed for the direct implementation of Twelf [20, 21]. Of special note here are optimizations like the linear heads treatment of unification described by Pientka and Pfenning [21] for minimizing occurs checking, that could make a difference in examples such as the *perm* program considered in the previous section: direct compilation would allow us to regain opportunities for such improvements that might be lost by translating first to λ Prolog and then relying on its implementation that is not specially optimized to treat Twelf-specific programs.

A more ambitious line of development concerns meta-reasoning over specifications. Existing tools might be used to reason about LF programs via the translation, the transparency of the translation becoming essential. Anecdotal evidence suggests that this transparency is not only enabling, it is also elucidating: that the generated *hohh* program is easier to reason about because it highlights those types that could have logical importance, and elides those that do not.

8. Acknowledgements

This work has been supported by the NSF grants CCR-0429572 and CCF-0917140. Opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] D. Baelde. *A linear approach to the proof-theory of least and greatest fixed points*. PhD thesis, Ecole Polytechnique, Dec. 2008. URL <http://www.lix.polytechnique.fr/~dbaelde/thesis/>.
- [2] D. Baelde, D. Miller, and Z. Snow. Focused inductive theorem proving. In J. Giesl and R. Haehnle, editors, *IJCAR*, Lecture Notes in Computer Science. Springer-Verlag, 2010. (to appear).
- [3] P. Brisset and O. Ridoux. The architecture of an implementation of lambda-prolog: Prolog/mali. In *ILPS Workshop: Implementation Techniques for Logic Programming Languages*, 1994.
- [4] A. Church. A formulation of the simple theory of types. *J. of Symbolic Logic*, 5:56–68, 1940.
- [5] A. Felty. *Specifying and Implementing Theorem Provers in a Higher-Order Logic Programming Language*. PhD thesis, University of Pennsylvania, Aug. 1989.
- [6] A. Felty and D. Miller. Encoding a dependent-type λ -calculus in a logic programming language. In M. Stickel, editor, *Proceedings of the 1990 Conference on Automated Deduction*, volume 449 of *LNAI*, pages 221–235. Springer, 1990.
- [7] A. Gacek. The Abella interactive theorem prover (system description). In A. Armando, P. Baumgartner, and G. Dowek, editors, *Fourth International Joint Conference on Automated Reasoning*, volume 5195 of *LNCS*, pages 154–161. Springer, 2008. URL <http://arxiv.org/abs/0803.2305>.
- [8] A. Gacek. *A Framework for Specifying, Prototyping, and Reasoning about Computational Systems*. PhD thesis, University of Minnesota, 2009.
- [9] A. Gacek, S. Holte, G. Nadathur, X. Qi, and Z. Snow. The Teyjus system – version 2, Mar. 2008. Available from <http://teyjus.cs.umn.edu/>.
- [10] A. Gacek, D. Miller, and G. Nadathur. Combining generic judgments with recursive definitions. In F. Pfenning, editor, *23th Symp. on Logic in Computer Science*, pages 33–44. IEEE Computer Society Press, 2008.
- [11] R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- [12] W. A. Howard. The formulae-as-type notion of construction. In J. P. Seldin and R. Hindley, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus, and Formalism*, pages 479–490. Academic Press, New York, 1980.
- [13] D. Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *J. of Logic and Computation*, 1(4):497–536, 1991.
- [14] D. Miller and A. Tiu. A proof theory for generic judgments. *ACM Trans. on Computational Logic*, 6(4):749–783, Oct. 2005.
- [15] D. Miller, G. Nadathur, F. Pfenning, and A. Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157, 1991.
- [16] G. Nadathur and D. Miller. An Overview of λ Prolog. In *Fifth International Logic Programming Conference*, pages 810–827, Seattle, Aug. 1988. MIT Press.
- [17] G. C. Necula. Proof-carrying code. In *Conference Record of the 24th Symposium on Principles of Programming Languages 97*, pages 106–119, Paris, France, 1997. ACM Press.
- [18] T. Nipkow. Functional unification of higher-order patterns. In M. Vardi, editor, *Proc. 8th IEEE Symposium on Logic in Computer Science (LICS 1993)*, pages 64–74. IEEE, June 1993.
- [19] F. Pfenning and C. Schürmann. System description: Twelf — A meta-logical framework for deductive systems. In H. Ganzinger, editor, *16th Conference on Automated Deduction (CADE)*, number 1632 in *LNAI*, pages 202–206, Trento, 1999. Springer.
- [20] B. Pientka. Eliminating redundancy in higher-order unification: A lightweight approach. In U. Furbach and N. Shankar, editors, *IJCAR*, volume 4130 of *Lecture Notes in Computer Science*, pages 362–376. Springer, 2006. ISBN 3-540-37187-7.

- [21] B. Pientka and F. Pfenning. Optimizing higher-order pattern unification. In *19th International Conference on Automated Deduction*, pages 473–487. Springer-Verlag, 2003.
- [22] J. Reed. Redundancy elimination for LF. *Electron. Notes Theor. Comput. Sci.*, 199:89–106, 2008. ISSN 1571-0661. doi: <http://dx.doi.org/10.1016/j.entcs.2007.11.014>.
- [23] Z. Snow. Parinati. <http://www.cs.umn.edu/~snow/parinati>, 2010.
- [24] Z. Snow. Realizing the dependently typed λ -calculus. Master's thesis, University of Minnesota, 2010.
- [25] Y. Takayama. Extraction of redundancy-free programs from constructive natural deduction proofs. *Journal of Symbolic Computation*, 12(1):29–69, 1991.

A. Proofs of Theorems

A.1 Correctness of the simplified encoding (Theorem 1)

A.1.1 Completeness

We use induction on the derivation of $\Gamma \vdash M : A$ to build one for $\llbracket \Gamma \rrbracket \longrightarrow \llbracket M : A \rrbracket$. We proceed by case analysis on the canonical type A .

If A is of the form $\Pi x:B.A'$ then M must be of the form $\lambda x:B.M'$ and the LF derivation must end with an *abs-obj* rule, i.e., a rule of the form

$$\frac{\Gamma \vdash A' : \text{Type} \quad \Gamma, x : B \vdash M' : A'}{\Gamma \vdash (\lambda x:B.M') : (\Pi x:B.A')} \text{ abs-obj}$$

The induction hypothesis gives us a derivation for

$$\llbracket \Gamma, x : B \rrbracket \longrightarrow \llbracket M' : A' \rrbracket.$$

By applying the rules $\forall R$ and $\supset R$ to this, we get a derivation for $\llbracket \Gamma \rrbracket \longrightarrow \forall x. \llbracket x : B \rrbracket \supset \llbracket M' : A' \rrbracket$. The righthand side of this sequent is the expected goal:

$$\llbracket (\lambda x:B.M') : (\Pi x:B.A') \rrbracket = \forall x. \llbracket x : B \rrbracket \supset (\llbracket A' \rrbracket (\langle \lambda x:B.M' \rangle x)),$$

and $\langle M' \rangle = (\langle \lambda x:B.M' \rangle x)$ by virtue of η -conversion.

If A is a base type then M must be of the form $x N_1 \dots N_n$ and the canonical LF derivation must end with a chain of *app-obj* rules following a *var-obj* rule that reveals that

$$x : \Pi y_1:B_1 \dots \Pi y_n:B_n.A' \in \Gamma.$$

Moreover, A must be $A'[N_1/y_1, \dots, N_n/y_n]$ and, from looking at the right upper premise of the *app-obj* rules, there must be shorter derivations of

$$\Gamma \vdash N_i : B_i[N_1/x_1, \dots, N_{i-1}/x_{i-1}]$$

for $1 \leq i \leq n$. By the induction hypothesis we obtain derivations \mathcal{D}_i of $\llbracket \Gamma \rrbracket \longrightarrow \llbracket N_i : B_i[N_1/x_1, \dots, N_{i-1}/x_{i-1}] \rrbracket$. Further, $\llbracket \Gamma \rrbracket$ must contain

$$\forall y_1. (\llbracket B_1 \rrbracket y_1) \supset \dots \supset \forall y_n. (\llbracket B_n \rrbracket y_n) \supset \text{hastype } (x y_1 \dots y_n) \langle A' \rangle,$$

i.e., the encoding of $x : \Pi y_1:B_1 \dots \Pi y_n:B_n.A'$. By applying *backchain* on that clause, choosing $\langle N_i \rangle$ for y_i and using the derivations \mathcal{D}_i , we obtain a derivation of

$$\llbracket \Gamma \rrbracket \longrightarrow \text{hastype } (x \langle N_1 \rangle \dots \langle N_n \rangle) (\langle A' \rangle [\langle N_1 \rangle / y_1, \dots, \langle N_n \rangle / y_n]).$$

The right side of this sequent is precisely

$$\llbracket (x N_1 \dots N_n) : A'[N_1/y_1, \dots, N_n/y_n] \rrbracket.$$

A.1.2 Soundness

We prove the soundness direction by induction on the derivation of $\llbracket \Gamma \rrbracket \longrightarrow (\llbracket A \rrbracket M)$: assuming that $\Gamma \vdash A : \text{Type}$ has a derivation, we establish that $M = \langle M' \rangle$ for some canonical object M' and we build a derivation of $\Gamma \vdash M' : A$. A case analysis on the structure of the canonical type A will guide us.

If A is of the form $\Pi x:B.A'$ then the structure of $\llbracket A \rrbracket$ forces the *hohh* derivation to conclude as follows:

$$\frac{\llbracket \Gamma, x : B \rrbracket \longrightarrow (\llbracket A' \rrbracket (M x))}{\llbracket \Gamma \rrbracket \longrightarrow \forall x. (\llbracket B \rrbracket x) \supset (\llbracket A' \rrbracket (M x))} \forall R, \supset R$$

Since A is a valid *Type* under Γ , B must also be, and A' must be valid under $(\Gamma, x : B)$. We can thus apply the inductive hypothesis, and we obtain that $M x = \langle M' \rangle$ and that $\Gamma, x : B \vdash M' : A'$ is derivable for some canonical object M' . Since x does not occur free in M , we conclude that

$$M = (\lambda x. \langle M' \rangle) = \langle \lambda x:B.M' \rangle,$$

and we derive $\Gamma \vdash (\lambda x:B.M') : (\Pi x:B.A')$ using the *abs-obj* rule and our derivation of $\Gamma \vdash B : \text{Type}$.

Otherwise, A is a base type, and the derivation we are considering is that of $\llbracket \Gamma \rrbracket \longrightarrow \text{hastype } M \langle A \rangle$. This derivation must end in a *backchain* rule that uses some clause in $\llbracket \Gamma \rrbracket$ of the form

$$\forall y_1. (\llbracket B_1 \rrbracket y_1) \supset \dots \supset \forall y_n. (\llbracket B_n \rrbracket y_n) \supset \text{hastype } (x y_1 \dots y_n) \langle A' \rangle;$$

note that the variables y_1, \dots, y_{i-1} can appear in $\llbracket B_i \rrbracket$ here. Thus, for some *hohh* terms N_1, \dots, N_n ,

$$\langle A \rangle = \langle A' \rangle [N_1/y_1, \dots, N_n/y_n],$$

$M = (x N_1 \dots N_n)$, and, for each i such that $1 \leq i \leq n$, there is a shorter derivation of

$$\llbracket \Gamma \rrbracket \longrightarrow (\llbracket B_i \rrbracket y_i) [N_1/y_1, \dots, N_i/y_i],$$

i.e., of $\llbracket \Gamma \rrbracket \longrightarrow (\llbracket B_i \rrbracket [N_1/y_1, \dots, N_{i-1}/y_{i-1}] N_i)$. Further, we know that $x : \Pi y_1:B_1 \dots \Pi y_n:B_n.A' \in \Gamma$ for some x . We now claim that, for $1 \leq i \leq n$, $N_i = \langle N'_i \rangle$ for some canonical LF object N'_i and that $\Gamma \vdash N'_i : B_i[N'_1/y_1 \dots N'_{i-1}/y_{i-1}]$ has a derivation. If this claim is true, then, we can use the *var-obj* rule to derive $\Gamma \vdash x : \Pi y_1:B_1 \dots \Pi y_n:B_n.A'$ and follow this by a sequence of *app-obj* rule applications to prove

$$\Gamma \vdash (x N'_1 \dots N'_n) : A'[N'_1/y_1 \dots N'_n/y_n].$$

Now, evidently $M = \langle x N'_1 \dots N'_n \rangle$ and, since substitution permutes with encoding, $A = A'[N'_1/y_1, \dots, N'_n/y_n]$. Thus, the desired result would be proven.

It only remains to establish the claim. We actually strengthen it to include also the assertion that, for $1 \leq i \leq n$,

$$\Gamma \vdash B_i[N'_1/y_1 \dots N'_{i-1}/y_{i-1}] : \text{Type}$$

has a derivation. To prove it, we use an inner induction on i . Since Γ is a well-formed context, and $x : \Pi y_1:B_1 \dots \Pi y_n:B_n.A' \in \Gamma$, there must be a derivation of

$$\Gamma, x_1 : B_1, \dots, x_{i-1} : B_{i-1} \vdash B_i : \text{Type}$$

for $1 \leq i \leq n$. Using Proposition 1 and the induction hypothesis we see that there must be a derivation of

$$\Gamma \vdash B_i[N'_1/y_1 \dots N'_{i-1}/y_{i-1}] : \text{Type}.$$

Noting that

$$\llbracket B_i \rrbracket [N_1/y_1, \dots, N_{i-1}/y_{i-1}] = \llbracket B_i[N_1/y_1, \dots, N_{i-1}/y_{i-1}] \rrbracket,$$

the outer induction hypothesis and the shorter derivation of

$$\llbracket \Gamma \rrbracket \longrightarrow (\llbracket B_i \rrbracket [N_1/y_1, \dots, N_{i-1}/y_{i-1}] N_i)$$

allows us to conclude that $N_i = \langle N'_i \rangle$ for some canonical LF term N'_i and that there is a derivation of

$$\Gamma \vdash N'_i : B_i[N'_1/y_1 \dots N'_{i-1}/y_{i-1}],$$

thus verifying the claim.

A.2 Completeness of the optimized encoding (Theorem 2)

If $\llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket M$ has a derivation, then $\llbracket \Gamma \rrbracket^+ \longrightarrow \llbracket A \rrbracket^- M$ has a derivation as well. Note that for this direction of the proof we are simply dropping information (subderivations) and so we do not rely on Γ being a valid specification or A being a valid type. We proceed by induction on the structure of the derivation of $\llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket M$, followed by case analysis on A .

If A is of the form $\Pi x:B.A'$ our derivation ends as follows:

$$\frac{\llbracket \Gamma, x : B \rrbracket \longrightarrow \llbracket A' \rrbracket (M x)}{\llbracket \Gamma \rrbracket \longrightarrow \llbracket \Pi x:B.A' \rrbracket M} \forall R, \supset R$$

By the inductive hypothesis $\llbracket \Gamma, x : B \rrbracket^+ \longrightarrow \llbracket A' \rrbracket^- (M x)$ has a derivation, and by applying $\forall R$ and $\supset R$ to this derivation we can construct a derivation of

$$\llbracket \Gamma \rrbracket^+ \longrightarrow \llbracket \Pi x:B.A' \rrbracket^- M$$

Otherwise, A is a base type and our derivation proceeds by backchaining on some $(y : \Pi x:\vec{B}.A') \in \Gamma$, with $\langle A \rangle = \langle A' \rangle[t_1/x_1 \dots t_n/x_n]$:

$$\frac{\llbracket \Gamma \rrbracket \longrightarrow F_1 \quad \dots \quad \llbracket \Gamma \rrbracket \longrightarrow F_n}{\llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket (y \vec{t})} \text{backchain}$$

Here, $F_i = (\llbracket B_i \rrbracket x_i)[t_1/x_1 \dots t_n/x_n]$. As in the completeness proof of the simplified encoding, we obtain by an inner induction that each t_i is of the form $\langle t'_i \rangle$ and thus that

$$F_i = \llbracket B_i[t'_1/x_1 \dots t'_n/x_n] \rrbracket (t_i).$$

We shall build the derivation of $\llbracket \Gamma \rrbracket^+ \longrightarrow \llbracket A \rrbracket^- (y \vec{t})$ by using *backchain* on the optimized encoding of $(y : \Pi x:\vec{B}.A') \in \Gamma$, by choosing \vec{t} for \vec{x} . The resulting premises are either

$$\llbracket \Gamma \rrbracket^+ \longrightarrow \llbracket B_i[t'_1/x_1 \dots t'_n/x_n] \rrbracket^- t_i$$

when x_i does not occur rigidly in A' , and this case is provided for by the inductive hypothesis, or \top otherwise, which we derive using $\top R$.